

AD-A132 522

ADVANCED HOUGH TRANSFORM IMPLEMENTATIONS(U) ROCHESTER
UNIV NY DEPT OF COMPUTER SCIENCE C M BROWN ET AL.
MAR 83 TR-113 N00014-82-K-0193


1/1

UNCLASSIFIED

F/G 12/1

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

12

ADVANCED HOUGH TRANSFORM
IMPLEMENTATIONS

C.M. Brown, M.B. Curtiss, and D.B. Sher
Computer Science Department
University of Rochester
Rochester, NY 14627

TR 113
March 1983

AD-A132522

DTIC FILE COPY

Rochester


Department of Computer Science
University of Rochester
Rochester, New York 14627

DTIC
ELECTE
SEP 15 1983

D

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By <u>Per Ltr. on file</u>	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	



ADVANCED HOUGH TRANSFORM IMPLEMENTATIONS

C.M. Brown, M.B. Curtiss, and D.B. Sher
Computer Science Department
University of Rochester
Rochester, NY 14627

TR 113
March 1983

Abstract

In the Hough transform, a feature in an input space I votes for parameters in a transform space P with which it is compatible. If it also uses negative votes to vote against parameters with which it is incompatible, background bias and inherent noise in P can be reduced and peaks sharpened.

In the usual Hough transform implementation, an array is used to accumulate votes, raising the problem that the space requirement increases exponentially with the number of parameters. Accumulation in a fixed-size content-addressable cache (or hash table) may provide a solution.

The preparation of this paper was supported in part by the Office of Naval Research under Grants N00014-80-C-0197 and N00014-82-K-0193, and the Defense Advanced Research Projects Agency under Grant N00014-78-C-0164.

Page - 1 -

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DTIC
ELECTED
SEP 15 1983

1. Definitions, Problems, Solutions

1.1 Hough Transform as Voting and Imaging

Considered as a voting process, the *Hough transform* maps features of a phenomenon in an input space I to sets of votes in the parameter space P . The parameter vectors receiving the most votes describe *instances* of the phenomenon explaining the most evidence. Ideally, the Hough transform maps a phenomenon into a set of votes that is a delta function in parameter space: the votes are unanimous. Usually the votes are more or less spread out, since several phenomena might generate a feature. Call the pattern of votes arising from a feature its *feature point spread function* (fpsf), and the superposition of all fpsfs from a phenomenon its *parameter point spread function* (ppsf) [Brown 1982]. A ppsf usually consists of a peak of votes at the correct parameter vector surrounded by *sidelobes* of incorrect votes (Fig. 1a).

1.2 Inherent Noise and Bias? Complementary Hough

Problem: When the input space contains multiple instances close together, sidelobes from one instance can overlap the peak from another, adding both variation and bias to the peak height. Even in single-instance noiseless inputs, the near-peak votes can spread the peak. **Solution:** In the *Complementary Hough transform*, the fpsf contains (negative) votes *against* parameters with which the feature is inconsistent, as well as (positive) votes *for* parameters with which it is consistent. The negative votes can cancel off-peak positive votes (Fig. 1b). Complementary Hough sharpens peaks and reduces sidelobe bias and variance.

1.2 Space Efficiency? Cache Hough

Problem: Multi-dimensional quantized parameter spaces are most straightforwardly implemented as arrays which grow impractically large as their dimension (the number of parameters) increases. **Solutions:** One solution is the use of dynamically quantized parameter space [O'Rourke 1981; Sloan 1981]. The solution explored here is *Cache Hough*: Accumulate vote tallies in a content-addressable cache (in hardware) or hash table (in software) [Brown and Sher 1982a; 1982b]. The abstract problem is that of estimating the mode(s) of a distribution using fixed memory. Cache Hough is a promising technique suitable for hardware implementation.

3. Results: Synthetic Data

3.1 Complementary Hough Performance

What fpsf for complementary Hough maximizes peak detectability and minimizes bias? All the results of this section arise from synthesized versions of circle detection for a constant radius, but the results are suggestive of general principles. In every case, the circumferential feature points in the input "image" contain no directional information; it is as if they were generated by a rotationally symmetric feature detector. This means that the fpsf in the regular Hough parameter space of circle centers is a complete circle of unit weight votes (each edge feature

votes for a set of centers at all locations a radius distance away). The fpsfs for complementary Hough are circles of unit-weight votes with neighboring concentric circles of negative votes. The first graph (Fig. 2) documents the search for the best complementary Hough fpsf. The data points correspond to fpsfs with the indicated negative surrounds. The fpsf with a circle of ones surrounded on both sides by $-1/2$'s was used for all the remaining work reported in this section.

The graphs of Figs. 3 - 9 document regular Hough and complementary Hough performance in seven different imaging situations, getting progressively more complex and noisy. As expected, complementary Hough can keep the mean background near zero, with small biases magnified by the contributions of many noise features (Fig. 3). Ideally, peak height is proportional to amount of evidence; complementary Hough achieve this (Fig. 4). Three signal-to-noise ratios are shown (Figs. 5 - 7), with complementary Hough comparing well in all three. SNR2 (Fig. 6) is meant to indicate the detectability of the peak by usual histogram peak-finders, while SNR3 indicates detectability by global thresholding (it is a multiple-instance version of the measure used below to quantify performance of cache Hough (Section 3.2)). Figs. 8 and 9 also show how reliably the peaks at true parameter vectors are distinguished from background by simple threshold detection.

3.2 Cache Hough Performance

We created synthetic phenomenon instances of controllable size, number, and location, embedded them in controllable amounts of noise features, and scanned the resulting images in several ways. Each time a feature was encountered, a vote burst (fpsf) of controllable size was emitted sequentially into a cache whose length and cache-flushing strategy were also controllable. The hope is that the parameter vector corresponding to the instance survives in the cache and accumulates the most votes. For technical details see [Brown and Sher 1982b]. Comprehensive results appear in [Brown and Sher 1982a].

Fig. 10 shows two synthetic images. Fig. 11 shows the arrangement of fpsfs and the resulting ppsfs in parameter space. Their form is dictated by theory in [Brown 1982]. The size of the instances and psfs and number of noise features in the image determine how many votes are sent to the cache. The scanning methods determine the order of sending. The cache length and flushing algorithm determine what vote tallies survive. We used four cache lengths (32, 64, 128, infinite). The infinite cache is equivalent to an accumulator array. Our instances had horizontal extent of 1, 2, or 3 and vertical extent of 11. The fpsfs were of 3, 9, and 15 votes. There was only one instance per image. Regular and Complementary Hough were used. Images were scanned five ways (Left to Right within Top to Bottom, Top to Bottom within Left to Right, Random with Replacement, Random without Replacement, and Left to Right within Random Scanline). One flushing method eliminated tallies below a threshold, the other eliminated a fixed fraction of tallies, chosen at random, below a larger threshold.

The results of these experiments yielded statistics such as the histograms of Figs. 12 and 13, which display a ratio

$$\text{SNR3} = \frac{(\text{votes for parameter vector describing instance})}{\max (\text{votes for other vectors in cache})}$$

If $\text{SNR3} > 1$, the correct parameter vector tallied a plurality of votes and would be detected as the largest entry in the cache. Otherwise, an incorrect parameter tallied at least as many (see also Fig. 7).

4. Results: Real Data

Our Hough Transform software laboratory allows flexible experimentation and comparison between methods. The cache accumulator module can be parameterized to provide finite or "infinite" length and various flushing methods. Sample output is shown in Fig. 14 and Table 1.

5. Discussion

Complementary Hough reduces inherent noise and bias and sharpens peaks at the cost of increased voting effort, depending on the implementation. In a naive implementation the cost might be strictly proportional to the number of votes, but in a connectionist implementation with hard-wired voting patterns there is no extra cost [Feldman and Ballard 1982]. Antialiasing is of some benefit, but also may affect voting cost.

Complementary Hough is most effective with dense voting that maximizes the superimposition of cancelling votes. If evidence is sparse, fpsfs and votes are sparse in the accumulator and little cancellation takes place; similarly if fpsfs are sparse. In the example of Section 4, positive votes occur in lines through 3-space, and each positive vote is surrounded by only six (of a possible 24 in 3-D) nearest-neighbor negative votes, resulting in less than optimal cancellation. If each feature votes for exactly one parameter vector (the fpsf is a single point), then complementary Hough degenerates to a linear enhancement operation: a central positive vote surrounded by a cloud of complementary ones is like a Laplacian operator.

In complementary Hough, sidelobes should cancel while the central peak accumulates positive votes. Feature perturbation error arises from quantization effects [Shapiro and Ianinno 1979]; e.g., the angle computed for an edge element is only approximate. Such error spreads the locus of the peak and superimposes complementary votes on it, degrading complementary Hough performance. Table 1 shows substantial feature perturbation in the angle of edge elements. Thus the results of Section 4 represent a severe test for the complementary Hough, and it is gratifying that it performs slightly better than regular Hough.

The remainder of the discussion concerns cache Hough. Finite caches of adequate size seem to be practical Hough accumulators, not significantly worse than accumulator arrays. The tables and histograms show a graceful and qualitatively predictable degradation of performance as noise and fpsf length increase and cache length decreases. Cache lengths of a few percent of the total number of votes seem adequate, but we do not yet know how this figure grows with parameter dimensionality.

The random (with replacement) scan was uniformly significantly worse in all cases than other methods. The random (without replacement) does not seem significantly better than nonrandom scans. Even in 30% noise, the Top to Bottom within Left to Right scanning method was not better than the Left to Right within Top to Bottom. This indicates (with our vertical instances) that successive bursts of votes (fpsfs) containing votes for the "true parameters" are not enough to help the peak establish itself. Probably the noise that arrives before the instance and the flush algorithm dominate the peak-establishment process. Thus the results are fairly insensitive to scanning method, as long as each input space point is considered once.

In good conditions (short fpsf, low noise, large instances), average cache complementary Hough performance is reliably markedly better than cache regular Hough. However, as conditions get worse, their average performance is approximately the same. Several of the histograms show a significant bimodality with one peak at 0. This indicates that the "true peak" never gets established in a significant number of cases, and that some pre-processing might help dramatically in establishing peaks. One idea is "pair Houghing," in which only parameter vectors that appear in two fpsfs are allowed into the cache. Clearly without such preprocessing the issue is whether peaks are able to survive in the cache, and this depends to a large extent on the Flush algorithm. The Random Below Threshold Flush makes for a significant improvement in (mean peak)/(maximum background) ratio, and reduces the frequency of occasions on which the correct parameters end up with zero vote weight.

Systematic distortion, occlusion, and feature perturbation pose thorny problems for cache Hough. An important topic of current work is to assess the performance of cache Hough schemes in more stringent noise conditions and to develop means of improving such performance.

Hierarchical caches can incorporate a form of "geometric locality" that is lost in simple content-addressable caches. Natural flush algorithms for a single-level cache are based strictly on tallies. Multiple caches in a resolution pyramid, maintained in parallel, can be used to restrict flushing to coherent localities in the input space. This idea is another topic of current research.

Figure Captions

Figure 1. (a) Plot of the ppsf for regular Hough circle location. Circle radius = 22, peak height = 156, average sidelobe = 3.88, sidelobe variance = 9.84. (b) Ppsf arising from complementary fpsfs consisting of a ring of 1's surrounded both inside and out by a ring of $-\frac{1}{2}$'s. Peak height = 156, average sidelobe = .26, sidelobe variance = 4.41.

Figure 2. The numbers beside the data points are the mean sidelobe (off-peak value) of the ppsf. The ordinate, called here SNR2, is a signal/noise ratio defined by the local peak height (its absolute height minus the height of its nearest neighbors) divided by the standard deviation of the sidelobe (non-peak) heights. The regular Hough voting pattern for circle detection yields the leftmost data point. Its local vote weight cross-section (a single 1) is shown below it. The "jagged" and "smooth" data points refer to the idea of anti-aliasing the fpsfs (in the computer graphics sense). In smooth voting, votes are put in a cell in amount proportional to the amount of fpsf passing through the cell. In regular Hough, this multiplies the number of votes and makes for a higher peak and proportionally higher mean background. The other data points represent more complicated fpsfs. The fourth, for example, is a circle of 1's and an inner circle of -1's. The second from the right is a circle of 1's surrounded by circles of $-1/2$'s.

Figure 3. Mean background. Figs 3-9 show complementary Hough performance in seven synthetic imaging situations, getting progressively more complex and noisy. Six ensembles, each of ten synthetic images, were analysed to yield each data point but the first, which is a baseline case arising from a single-instance (noiseless) image. In each of the images in an ensemble, five shape instances and some noise features appear. The images in an ensemble randomly differed in the placement of the instances and noise points. The six ensembles systematically differed in the percentage of the instance features visible in the image (100, 75, 50, and 25%) and in the percentage of noise features in the image (0, 15, and 30% of the image features). The last two ensembles are similar except the multiple features were closer together in the "Dense" ensemble. Where they appear, variation bars encompass one standard deviation and are representative of other data point variations.

Figure 4. Fidelity. This is measured by the peak height normalized by how much image evidence contributes to it. Ideally this measure stays constant.

Figure 5. SNR1. The raw (absolute) peak height divided by the standard deviation of off-peak heights.

Figure 6. SNR2. SNR2 is the local peak height (its height above its neighbors) divided by the standard deviation of off-peak heights. SNR2 is often a better measure of peak detectability than SNR1.

Figure 7. SNR3. SNR3 measures peak detectability when simple thresholding is the detection criterion. SNR3 is the average peak height divided by the maximum off-peak height. This is a multiple-instance version of the measure used for the cache-based Hough study of Section 3.2.

Figure 8. No False Negatives. This shows how many parameter vectors received at least as many votes as the true peaks. Equivalently, how many wrong peaks are detected by a simple thresholding strategy that sets its threshold just low enough to include all the true peaks.

Figure 9. No False Positives. This shows how many true peaks have more votes than all non-peak vectors. Equivalently, how many of the true peaks are detected by a simple thresholding strategy that sets its threshold just high enough to exclude all non-peak vectors.

Figure 10. (a) Two 18-feature, two-part instances in a 20x20 abstract image. (b) The instances of (a) with 15% signal-independent noise.

Figure 11. (a) The features produce straight lines of votes in parameter space yielding a rosette-like ppsf, as suggested by an analytic model [Brown 1982]. For complementary Hough, lines of positive votes (weight 2) are flanked by negative votes (weight -1). (b) Ppsfs for 13-feature regular Hough and Complementary Hough.

Figure 12. Performance varying with cache length and complementary Hough. For regular Hough, the mean number of votes per run was 639 and the histogram records 100 runs. For complementary Hough, there were twenty runs and the mean number of votes was 1917.

Figure 13. Performance varying with flushing algorithm. With 30% noise, the mean number of votes cast per run is 1179.

Figure 14. (a) An input image. (b) A projection of the 3-D Hough and complementary Hough accumulator arrays with parameters of circle location and radius. (c) The ten circles with the most votes with regular Hough. (d) As in (c) with complementary Hough.

TABLE 1

Statistics for NORMAL HOUGH:
Results for circles in order

Rank	Radius	CenterX	CenterY	Angle Diff	Support
0	14	154	97	0.22	220.05
1	14	115	98	0.16	216.92
2	14	114	58	0.16	238.09
3	14	153	58	0.24	210.71
4	11	154	100	0.27	193.63
5	17	155	94	0.19	184.54
6	14	159	235	0.15	262.88
7	11	113	100	0.34	134.79
8	8	114	64	0.53	126.77
9	17	115	95	0.20	164.39
Avg:				0.25	195.28


Statistics for COMPLEMENT HOUGH:
Results for circles in order

Rank	Radius	CenterX	CenterY	Angle Diff	Support
0	14	154	97	0.22	220.05
1	14	115	98	0.16	216.92
2	14	114	58	0.16	238.09
3	11	154	100	0.27	193.63
4	14	153	58	0.24	210.71
5	17	155	94	0.19	184.54
6	14	159	235	0.15	262.88
7	8	114	64	0.53	126.77
8	11	113	100	0.34	134.79
9	14	120	235	0.21	190.24
Avg:				0.25	197.86


TABLE 1. The radii, location, and rank ordering by votes of multiple circle instances. Also their support in the image, measured by comparing the expected edges in the instance to the real edges found by a Sobel edge detector in the image at the corresponding points. Angle difference is the average angle between the expected and actual edges, weighted by actual edge strength. Support is the average edge strength for edges within $\pi/10$ of the expected angle.

References

- Ballard, D.H., "Generalizing the Hough transform to detect arbitrary shapes," *Pattern Recognition* 13, 2, 111-122, 1981.
- Brown, C.M., "Bias and noise in the Hough transform I: Theory," TR 105, Computer Science Dept., U. Rochester, June 1982.
- Brown, C.M. and D.B. Sher, "Hough transformation into cache accumulators: Considerations and simulations," TR 114, Computer Science Dept., U. Rochester, August 1982a.
- Brown, C.M. and D.B. Sher, "Modeling the sequential behavior of Hough transform schemes," *Proc., DARPA Image Understanding Workshop*, 115-123, Palo Alto, CA, September 1982b.
- Duda, R.O. and P.E. Hart, "Use of the Hough transform to detect lines and curves in pictures," *CACM* 15, 1, 11-15, 1972.
- Feldman, J.A. and D.H. Ballard, "Connectionist models and their properties," *Cognitive Science* 6, 205-254, 1982.
- Hough, P.V.C., *Method and Means for Recognizing Complex Patterns*, U.S. Patent 3069654, 1962.
- O'Rourke, J., "Dynamically quantized spaces for focusing the Hough transform," *Proc., 7th IJCAI*, 737-739, Vancouver, B.C., August 1981.
- Shafer, S.A. and T. Kanade, "Using shadows in finding surface orientations," *Proc., DARPA Image Understanding Workshop*, 90-102, Palo Alto, CA, September 1982.
- Shapiro, S.D. and Iannino, A., "Geometric constructions for predicting Hough transform performance," *IEEE Trans. PAMI-1*, 3, July 1979.
- Sloan, K.R., Jr., "Dynamically quantized pyramids," *Proc., 7th IJCAI*, 734-736, Vancouver, B.C., August 1981.



(a)



(b)

Figure 1.

$SNR2 = (PEAK - AVE(WINGS)) / STDEV$
 RADIUS = 40
 ▲ -- JAGGED
 ● -- SMOOTH

PEAK = 304
 PEAK = 3596

SNR2
 150 -





Figure 2. NOISELESS PARAMETER POINT SPREAD FUNCTION

$SNR1 = PEAK/STDEV$
 R = 10
 ● -- REG. NOISE (1)
 ■ -- CHOUGH (-.5, 1, -.5)




Figure 5.

PEAK BACKGROUND

R = 10
 ● -- REG. NOISE (1)
 ■ -- CHOUGH (-.5, 1, -.5)

AVE(OFF-PEAK)
 20 -





Figure 3.

INFIDELITY

R = 10
 ● -- REG. NOISE (1)
 ■ -- CHOUGH (-.5, 1, -.5)

PEAK WEIGHT
 INSTANCE FRACT.
 200 -



Figure 4.




Figure 6.



Figure 7.






Figure 8.

Figure 9.




(a)

Figure 10.




(b)




(a)

Figure 11.

(b)




(c)



Cache length: 32. Scan: ltoB, ltoH.
 Noise: 15%. Instances: 1.
 Parts: 1. fspf length: 9.
 Part length: 11. N: 100

Flush: threshold

Mean of Ratios: 1.212250
 Std. Dev. of Ratios: 0.972063




(a)

Cache length: 64. Scan: ltoB, ltoR.
 Noise: 15%. Instances: 1.
 Parts: 1. fspf length: 9.
 Part length: 11. N: 100

Flush: threshold

Mean of Ratios: 1.805774
 Std. Dev. of Ratios: 0.501593




(b)

Figure 12.

Cache length: 64.
 Scan: ltoB, ltoH.
 Noise: 15%

fsf length: 9




(c)

Cache length: 64. Scan: ltoB, ltoR.
 Noise: 30%. Instances: 1.
 Parts: 1. fspf length: 9.
 Part length: 11. N: 100

Flush: threshold


Mean of Ratios: 0.895691
 Std. Dev. of Ratios: 0.645508



(a)


Flush: Random Below Threshold

Mean of Ratios: 0.893294
 Std. Dev. of Ratios: 0.611755




(b)


Figure 13.




(a)



(b)



(c)



(d)

Figure 14.

END

FILMED

9-83

DTIC